

BEST AVAILABLE COPYAttorney Docket No. 35399.18
Customer No. 000027683**REMARKS**

Claims 1-22 were pending in the application. Claims 1, 2, 4, 9, 12-14, 16, 18 and 19 were rejected. Claims 3, 5-8, 10, 11, 15, 17 and 20-22 were objected to. Claims 1-22 remain active in the application. In view of the following remarks, reconsideration of the application is respectfully requested.

Response to Drawing Objection

Figure 1 was objected to "because item 12 should be labeled with 'network processing device' (37 CFR 1.83(a))." (Office Action, paragraph 1.) Regrettably, Applicant does not understand the basis for the objection, and therefore respectfully traverses. Applicant's copy of Figure 1 contains a dashed box, clearly labeled "12". The specification states that "FIG. 1 shows a network processing device 12 connected to an Internet network 14." (Page 3, ll. 6-7.) Thus item 12 is already identified as a network processing device. Therefore, Figure 1 meets the requirement of 37 C.F.R. § 1.83(a) that "the drawing in a nonprovisional application must show every feature of the invention specified in the claims." Applicant respectfully requests that the objection to the drawings be withdrawn.

Response to Claim Rejections***Claims 1, 2, 4, and 9***

Claims 1, 2, 4, and 9 were rejected under 35 U.S.C. § 102(e) as anticipated by U.S. Patent No. 6,810,031 ("Hegde"). Applicant respectfully traverses this rejection, and submits that Hegde fails to teach each element of any of these claims, and therefore does not anticipate claims 1, 2, 4, and 9.

The rejection states, in part, that Hegde teaches a data rate controller comprising "a bandwidth limiter to identify a maximum allowable bandwidth for an input port (col. 7, lines 48-55)..." (Office Action at 3, paragraph 2.) Applicant respectfully disagrees. The section of Hegde cited for support discusses a "Credit variable." The Credit variable allows each card to "bank" unused byte grants from previous cycles, up to a defined Credit Threshold. (Hegde, col. 7, ll. 50-58.) Each card is allowed to request, for a cycle, an additional bandwidth allocation that includes a "draw-down" less than or equal to its accumulated Credit. (*Id.*, col. 8, ll. 2-3.) The card requests for each cycle a bandwidth that is the sum of the draw-down and its "Accumulated Buffer Occupancy." (*Id.*, col. 8, ll. 8-19; col. 9, ll. 19-26.)

Comparing this teaching of Hegde to the limitation of claim 1, the two are quite different. Claim 1 requires "a bandwidth limiter configured to identify a maximum allowable bandwidth the

BEST AVAILABLE COPYAttorney Docket No. 35399.18
Customer No. 000027683

input port is allocated on the backplane." Hegde's Credit and Credit Threshold fail to function as the claimed bandwidth limiter limitation. The Credit and Credit Threshold merely limit the maximum size of a requestable "draw-down", but they fail to limit allocated bandwidth, or even requested bandwidth, which is a non-limited function of Accumulated Buffer Occupancy. Thus the Credit counter and associated draw-down mechanism cannot identify a maximum allowable bandwidth, as claimed.

The rejection further states, in part, that Hegde teaches a data rate controller comprising "a bandwidth tracker to identify an allocated bandwidth and to prevent the input port from connecting to the output port when the bandwidth is used up (col. 7, line 66 through col. 8, line 5.)" (Office Action at 3, paragraph 2.) This Hegde reference is also to the Credit variable and its use for requesting a draw-down. As explained above, the draw-down is an "extra" that can be added to the requested bandwidth on top of Accumulated Buffer Occupancy. (Hegde, col. 9, ll. 20-26.) Hegde's Credit and draw-down cannot function as the "bandwidth tracker" in claim 1, at least because it can never prevent the card from requesting (and receiving) bandwidth based on Accumulated Buffer Occupancy, no matter what actual bandwidth the card has been using.

Based at least on the differences identified above, Applicant respectfully submits that Hegde cannot anticipate claim 1 and its dependent claims 2, 4, and 9.

Regarding claim 2, Hegde also fails to teach "a register that stores a programmable peak time slot rate value." Hegde's Credit has nothing to do with peak time slot rate value as claimed. The Credit (and its threshold) are not based on rate at all, but represent aggregate bytes unused from previous grants. (Hegde, col. 7, ll. 50-55.)

Regarding claim 4, Hegde's Credit counter is decremented when the card "is allowed to transmit an extra amount of traffic to a card (draw-down)." (*Id.*, col. 8, ll. 1-5, emphasis added.) And the counter is incremented when, for a given cycle, the card "transmits fewer bytes [to the target] than assigned." (*Id.*, col. 7, ll. 52-54.) Thus the decrement does not occur "when the input port is connected through the backplane to an output port" as claimed, but only in the special case of an extra draw-down in addition to regular traffic. Likewise, the increment does not occur "when the input port is not connected through the backplane to the output port" as claimed, but only in the special case where a grant is made but not fully used.

Regarding claim 9, the rejection refers to Hegde's processors running a Bandwidth Distribution Protocol (BWDP) as teaching the claimed "arbitration circuit." Claim 9 states that the arbitration circuit is "configured to arbitrate between input ports for connections to output ports," and selects "the input ports for a next time slot according to both a priority and weight of packets at the input ports." Hegde's BWDP does not arbitrate for connections during a time

BEST AVAILABLE COPYAttorney Docket No. 35399.18
Customer No. 000027683

slot—his switch fabric is connectionless, relying on switch fabric buffers. (Hegde, col. 7, ll. 19-22.) The input cards are therefore not contending by arbitration for selection to a next time slot connection, but merely told to share an output port nicely during the next time slot: “[s]o, as it should be apparent, the BWDP algorithm of the present invention is necessary to **allocate bandwidth** across the switch fabric 100 fairly among the two contending IPE cards 104.” (*Id.*, col. 7, ll. 22-39, emphasis added.) Bandwidth sharing between multiple input cards within a cycle is not the same as selection for an input port-to-output port connection for a cycle.

Claims 12-14 and 16

Method claims 12-14 and 16 were rejected under 35 U.S.C. § 102(e) also, as anticipated by Hegde. Applicant respectfully traverses this rejection, and submits the Hegde fails to teach each element of any of these claims, and therefore does not anticipate claims 12-14 and 16.

Regarding claim 12, many of the observations noted above exemplify the distinctions between the limitations of claim 12 and Hegde. For instance, claim 12 requires “sending requests from the input ports for connecting to the output ports during a next time slot.” As noted, Hegde uses a connectionless fabric, and all cards that want bandwidth are given some bandwidth during a cycle. The cards request a desired amount of bandwidth for a cycle, but not a connection, which is unnecessary in this architecture. (Hegde, col. 9, ll. 19-29.)

Also, because Hegde is connectionless, it does not increase or decrease bandwidth allocation based on whether a connection is granted during a time slot, as claimed. As noted above, Hegde's Credit counter increases for unused bytes granted while an input card is allowed to transmit to an output card, and does not increase when no connection exists as claimed. And the Credit counter decreases when extra draw-down bytes are allowed during a cycle, not “for the input ports that are connected to the output ports” as claimed.

Further, Hegde does not teach “preventing the input ports from sending requests for the input ports when the bandwidth allocated to the input ports has been exhausted.” In fact, Hegde does not teach any condition under which a card can exhaust a bandwidth allocation. Each card is allowed to request, for each cycle, at least its Accumulated Buffer Occupancy for priority and non-priority transfers. (Hegde, col. 9, ll. 21-29.)

Based at least on the differences identified above, Applicant respectfully submits that Hegde cannot anticipate claim 12 and its dependent claims 13, 14, and 16.

Further regarding claim 13, Hegde fails to teach disabling bandwidth allocation to input ports that reach a maximum allowable bandwidth allocation as claimed. Although the rejection does not point to any particular passage of Hegde for support, Hegde's Credit Threshold does

BEST AVAILABLE COPYAttorney Docket No. 35399.18
Customer No. 000027683

not disable bandwidth allocation, but merely limits the maximum extra draw-down bytes that a card can request.

Claim 14 requires "providing a programmable amount of bandwidth allocation for the input ports." The rejection does not point to any teaching of Hegde for support. To Applicant's reading, Hegde apportions bandwidth according to buffer fullness, not according to a programmable amount of bandwidth allocation.

Regarding claim 16, Applicant again points out that Hegde uses a connectionless switch fabric. Thus Hegde does not use the claimed method of preventing output ports that have used up their allocated bandwidth from granting connections to the input ports. Instead, Hegde iterates from highest priority to lowest priority (col. 10, ll. 32-33), and apportions bandwidth grants at each priority among the requesters (col. 11, l. 45 to col. 12, l. 6, corresponding to Figure 11 noted by the Examiner) Although it is true that this algorithm will not **increase** the amount of its grants further for a cycle once all available bandwidth for the cycle has been granted, connections are never granted or not granted for an output, as the outputs always start with some bandwidth for each cycle, and this bandwidth is apportioned first to the highest priority requesters. In claim 16, on the other hand, an output port cannot grant a connection at all for a next time slot if it has used up an allocated bandwidth.

Claims 18 and 19

Claims 18 and 19 were rejected under 35 U.S.C. § 102(e) as anticipated by U.S. Patent Application No. 2004/0081185 ("Grow"). Applicant respectfully traverses this rejection, and submits the Grow fails to teach each element of either of these claims, and therefore does not anticipate claims 18 and 19.

Claim 18 requires "a set of data rate controllers associated with each one of the virtual output queues for controlling a data rate that the input ports can transfer data to the output ports over the switch fabric." Grow fails to provide at least this element of claim 18.

The rejection identifies a "frame selector" (item 16) in Figure 3 of Grow as corresponding to this element, and notes that the frame selector serves a set of queues by a scheduler that can use, e.g., weighted priority scheduling. The frame selector does not, however, function as a data rate controller, but merely determines a relative frequency via which each queue will be served, assuming that the queue has data waiting and that the queue's destination is not blocked (paragraphs 46-49). In fact, Grow has no need to control data rate through his switch fabric at all. Paragraph 18 teaches that each input port has the ability to submit frames to the crossbar 6 at **twice** the rate that it receives them, therefore the frame selector never has to even

BEST AVAILABLE COPYAttorney Docket No. 35399.18
Customer No. 000027683

worry about data rate for a queue. As Grow teaches, "buffering at the crossbar 6 using RAM in combination with the increased rate of transmission between the input ports and the crossbar 6 enables frames to be forwarded to the output ports 4 at a rate greater than the media speed (i.e., the data rate at which data frames are received at the input ports 2)." (Grow, para. 18, ll. 4-9.) Grow also teaches that if too many inputs target a single output, the switch fabric uses dropping algorithms (as opposed to rate control) to alleviate the condition. (*Id.*, para. 48, ll. 15-22.) Based on Grow's silence as to per-queue rate control, and his instigation of an "overengineered" pipe to the switch fabric and a dropping algorithm in the switch fabric, it cannot be assumed that Grow uses any type of rate control at the input queues.

Response to Claim Objections

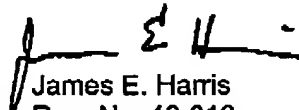
Claims 3, 5-8, 10-11, 15, 17, and 20-22 were objected to as being dependent upon a rejected base claim, but otherwise allowable. In view of the arguments presented above for the patentability of the base claims from which each objected-to claim respectively depends, Applicant has elected at this time to leave these as dependent claims.

Conclusion

Applicant respectfully requests that the rejections of claims 1, 2, 4, 9, 12-14, and 16 be withdrawn for the reasons presented above, and that the application be allowed in present form.

Respectfully submitted,

Date: 10/26/05
HAYNES AND BOONE, LLP
901 Main Street, Suite 3100
Dallas, Texas 75202-3789
Telephone: 512.867.8502
Facsimile: 512.867.8663
ipdocketing@haynesboone.com


James E. Harris
Reg. No. 40,013

CERTIFICATE OF FACSIMILE TRANSMISSION

I hereby certify that this paper is being facsimile transmitted to the Patent and Trademark Office on the date shown below.

Facsimile No.: 571-273-8300Date: 10/26/05


Susan C. Lien